

"Express Mail" mailing label number:

EL329124089US

FIFO WITH UNDO-PUSH CAPABILITY

Jennifer Pencis
Chandrakant Pandya
Sanjiv K. Lakhanpal
Mark D. Nicol

5

RELATED APPLICATIONS

This application is related to the following applications filed the same day as the present application, which applications are incorporated herein by reference in their entirety:

10 *Int* Application serial number _____ (Attorney Docket Number 1001-0129),
Al entitled "Command Receive Filter", naming Jennifer Pencis, Chandrakant Pandya,
Sanjiv K. Lakhanpal and Mark D. Nicol as inventors; and

Int Application serial number _____ (Attorney Docket Number 1001-0128),
Al entitled "Out of Order Probing in An In-Order System ", naming Sanjiv K.
15 Lakhanpal, Jennifer Pencis, Chandrakant Pandya, and Mark D. Nicol as inventors.

BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to commands received by an integrated circuit and more particularly to commands which have a possibility of being canceled.

20 Description of the Related Art

In current processor implementations, a processor issues requests to a memory controller for such transactions as, e.g., read or write commands, or probe responses. In one implementation for a stand-alone processor, a command channel is provided over which the processor issues requests to another integrated circuit incorporating
25 the memory controller. In such systems, data may be transmitted to and from the processor over a data channel that is separate from the command channel. In a typical

computer system incorporating such a processor, an integrated circuit interfaces to the processor to provide both a memory control function and a bridge function between the host bus (the command and data channel) and other system buses. One of the bridge functions typically provided is a bridge between the host bus and the industry
 5 standard Peripheral Component Interconnect (PCI) bus. Such integrated circuits have been described in the art as "north bridges".

One of the functions performed by memory controller logic on a north bridge is to probe the cache memory located on the processor in order to maintain coherence between the various components of the memory when another device, e.g., an I/O
 10 device, accesses memory. The probe determines whether the cache memory on the processor has a copy of the memory location (typically a cache line containing that memory location) about to be accessed by the I/O device. The processor responds through the command channel with a probe response indicating whether or not that particular cache line is located in the cache memory.

15 The command channel may be a split transaction bus in that commands such as reads may be separated from the data that is returned. Because multiple transactions may be outstanding, the processor tracks the number of outstanding transactions. The commands issued by the processor are received by the north bridge into a command queue.

20 The command channel may be relatively narrow, e.g., 13 bits, and therefore command/address information is sent over the command channel in multiple clock cycles. However, because certain commands are executed speculatively, a command may be canceled after a portion of the command has already been sent over the command channel. Therefore, it would be desirable to handle the cancellation of
 25 commands with a minimum effect on downstream logic.

SUMMARY OF THE INVENTION

Accordingly, the invention provides in one embodiment an integrated circuit that receives command information for one command over a plurality of bit times. The integrated circuit includes a command queue storing portions of command
 30 information received into the integrated circuit during consecutive bit times. The

integrated circuit further includes control logic responsive to a cancellation indication in the command information, which indicates that the current command is canceled, to repoint a write pointer to point to the canceled command already stored in the command queue.

5 In another embodiment, the invention provides a method for storing command information into a command queue in an integrated circuit. The method includes receiving a plurality of command segments comprising one command in a plurality of phases, each command segment being received in a different phase. The command segments are pushed into a command queue. Control logic checks for a cancellation
10 indication for the command being received. If a cancellation indication is received, the control logic for the command queue performs an undo-push operation to remove the command segments stored in the command queue associated with the cancelled command.

BRIEF DESCRIPTION OF THE DRAWINGS

15 The present invention may be better understood, and its numerous objects, features, and advantages made apparent to those skilled in the art by referencing the accompanying drawings in which the use of the same reference symbols in different drawings indicates similar or identical items.

20 Fig. 1 is a block diagram of a system in which a processor is coupled to a north bridge by a command and data channel.

Fig. 2 illustrates a receive circuit according to an embodiment of the present invention.

Figs. 3A and 3B are exemplary packets for non-probe response commands and probe response commands, respectively.

25 Fig. 4 is a flow chart illustrating operation of filtering commands received in the receive circuit.

Fig. 5 is a flow chart illustrating responding to a cancellation indication to perform an undo-push operation.

Fig. 6 is a flow chart illustrating operation of adjusting the content count of the command queue in response to various push, pop and undo-push combinations.

DESCRIPTION OF THE PREFERRED EMBODIMENT(S)

Referring to Fig. 1, the block diagram illustrates an exemplary computer system suitable for incorporating an embodiment of the invention. Processor 101 is coupled to north bridge 103 through unidirectional command channels 105 and 107 as well as a bi-directional data channel 109. The other control signals and information signals between the north bridge 103 and processor 101 are not shown to simplify the diagram. North bridge 103 provides a bridge function between the host bus (including the command and data channels) and the PCI bus 111 and Accelerated Graphics Port (AGP) 113. In addition, north bridge 103 provides a memory control function for system memory 115.

Referring to Fig. 2, north bridge 103 includes a receive circuit 201 coupled to receive commands from processor command channel 107. In one embodiment, the processor bus is a split transaction bus utilizing clock forwarding technology. In one embodiment data is forwarded on each edge of the clock. Thus, each edge of the clock is a bit time. Clock forwarding refers to supplying a clock along with data on the command channel. An exemplary bus of that type is the processor bus utilized by Digital Equipment Corporation's Alpha™ EV6 or by Advanced Micro Devices Athlon™ processor. The receive circuit translates the incoming signals to the local time domain. In addition, the receive circuit includes a receive filter to remove unused commands as described further herein.

While the command channel is in idle state no operations (NOPs) are present in the channel. In one embodiment, the command channel data, which typically includes command and address information, is transferred in four bit times, i.e., four clock edges. The command channel information may include, e.g., a bit field defining the command as a read command and bit fields defining the start address and length of the read operation. The first portion of a command is sent in bit time 0. The bit time is counted in a bit time counter (not shown) that tracks the current bit time. In the illustrated embodiment, the bit time counter counts bit times from 0 to 3. As seen in Fig. 2, data from command channel 107 is clocked into one of four sets of buffers,

203-206. In one embodiment, the buffers are formed of positive edge and negative edge flip-flops to account for the fact that data is forwarded on the command channel on each clock edge. Each set of buffers in the embodiment illustrated includes 13 flip-flops so each buffer can store thirteen bits of command channel information. The command and address information from the first bit time is stored in buffer 203, the second bit time in buffer 204, the third bit time in buffer 205 and the fourth bit time in buffer 206. The enable bits [3:0] selectively enable the buffers 203-206 according to the bit time determined by the bit time counter.

A multiplexer select counter 207 provides multiplexer selects that are used to unload the command information from the receive buffers through multiplexer 208. The multiplexer select can be set to one of four values, and determines the sequence for unloading the buffers, thus accounting for differences in time domains between the command channel and the receive circuit. Once the buffers are unloaded, the command information can be assembled into the 52 bit command information in a command storage queue 212 utilized in the illustrated embodiment.

First in first out buffers (FIFOs) 213, 215, 217, and 219 provide a command storage queue 212 for a command interpreter 224. The command storage queue includes four sixteen deep 13 bit FIFOs which can assemble sixteen 52 bits commands. However, as previously mentioned, not all of the commands received on command channel 107 require a response from north bridge 103. Accordingly, a command filter 220 is included to filter out those commands which do not require a response. Those commands not requiring a response are identified in filter 220, either as hard coded values for comparison or programmable values to provide greater flexibility. The command information received from command channel 207 is compared to the predetermined commands that do not require a response. Filter 203 may be implemented as series of comparators to compare, e.g., a five bit command field in the command information to the predetermined group of commands that are not forwarded to the command queue. If the received command is not in the predetermined group of commands, then the command is forwarded to command storage queue 212. However, if the received command is one of the predetermined commands, then the received command is dropped. That way, the command interpreter does not have to deal with those commands that do not require a response.

In one embodiment the commands not requiring a response from the north bridge include NOP (used for idle cycles on the command channel), a NZNOP (containing probe information), a victim data buffer flush (VDBFlush) (related to maintaining cache coherency for cache lines evicted from the cache memory), and probe responses. Those commands are filtered out at filter 220. However, the probe response command is not simply dropped as are the other commands not requiring a response from north bridge 103.

Exemplary command packets received and filtered by filter 220 are shown in Figs. 3A and 3B. The illustrated command packets take four bit times to transfer with the five bit command field being transferred in bit time 0. NOP commands occupy two bit times rather than four and thus can start in bit time 0 or 2. The PA fields provide the physical address bits corresponding to the command. The mask field MASK[7:0] is used for sub-cache block commands (byte, longword, or quadword) and indicates which bytes, longwords, or quadwords on data channel 109 are valid for the requested data transfer. The M1, M2 and CH bits relate to probe responses and their use is described further herein. The ID[2:0] field specifies a buffer entry in the processor corresponding to the command in CMD[4:0].

The command packet illustrated in Fig. 3B is used for certain types of probe responses. As previously mentioned, the probe responses are not simply dropped as are the other commands not requiring a response from north bridge 103. Instead, in order to avoid deadlock, probe responses should be processed out-of-order with respect to other commands in command queue 212. The probe response is forwarded to probe response logic 222. Depending on the nature of the probe response, information may be provided to command interpreter through signal lines 221 to indicate to the command interpreter that a data movement command or other command, needs to be sent to the processor. Signal lines 221 may convey a victim data buffer (VDB) buffer location in the processor. In other cases, the probe response information may require no action on the part of the command interpreter. For example, the probe response information indicate a probe miss and that information may be conveyed to allow, e.g., a DMA operation to complete.

There can be several types of probe responses. Probe responses may occur without data movement indicating a cache line hit or miss that does not require data movement. Those probe responses can be embedded in other commands (see Fig. 3A) using the M1, M2 and CH bits. Probe responses may also occur with data movement, e.g., when the cache line is "dirty", meaning it has been modified and differs from the copy of the cache line in system memory. If a cache line is dirty, it needs to be written back into system memory. Probe responses may also be received indicating that there is a command in-flight affecting the cache line that is the subject of the probe response. That means that the probe response has hit in, e.g., the victim data buffer.

A probe response not requiring data movement may be included in the command packet illustrated in Fig. 3A. Several bits (M1, M2, CH) are used to encode the probe status of the oldest outstanding probe issued by the system to the processor. In one embodiment, those bits are encoded as shown in Table 1:

TABLE 1

M1	M2	CH	Description
0	0	0	No probe response
0	0	1	Invalid
0	1	0	Cache miss on oldest outstanding probe (bit-time 2)
0	1	1	Cache hit on oldest outstanding probe (bit-time 2)
1	0	0	Cache miss on oldest outstanding probe (bit-time 0)
1	0	1	Invalid
1	1	0	Invalid
1	1	1	Invalid

Those bits are checked by the command filter and the information forwarded so that the probe response can be disposed of by the system out-of-order with respect to commands in the command queue. The command itself containing the encoded probe status is entered into the command queue for normal in-order processing.

When the command is processed, those bits relevant to the probe response can be ignored. A probe response requiring data movement or a probe response that hits in the victim data buffer (VDB) in the processor is provided in a separate command packet (see Fig. 3B) dedicated to probe responses.

Sub 7
A3

The probe response logic 222 may be sufficient to hold multiple probe responses or in a simple implementation, only a single probe response. The probe response logic processes the probe responses both from the probe response commands and extracted from other commands, out-of-order with respect to command queue 212 to avoid possible deadlock situations.. Further information on handling of probe response information can be found in application serial number _____ (Attorney Docket Number 1001-0128), entitled " Out of Order Probing in An In-Order System", previously incorporated by reference.

Referring to Fig. 4, the flow diagram illustrates the operational flow of the receive circuit and related logic. In 401, a first portion of the 52 bit command information, received in bit time 1, is received into buffer 203. Filter circuit 220 receives the contents of buffer 203, which is selected by muxselect counter 207. In 403, filter circuit 220 determines if the command field in the first portion of the command channel information matches one of the predetermined commands. If not, then the command information is stored in command queue 212 in 405. If the command is one of the predetermined commands, then the receive filter determines in 407 whether or not the command is a probe response. If so, the probe response is supplied to probe response logic 222 in 409. Otherwise, the commands are dropped in 411. That means that all command segments associated with the command segment checked in filter 220 are dropped and not stored in command queue 212. The receive filter then continues receiving additional command information. Although not illustrated in Fig. 4, encoded probe response information present in non-probe response commands is extracted by the command filter and is also forwarded to probe response logic 222 for out-of-order processing.

While the above description assumed the first portion of the 52 bit command information included the relevant command information to compare, one of skill in the art will appreciate that the relevant command information may be received in any phase. In fact, while filter 220 has been described as performing a filter operation based on a particular command, in fact the filter may check for any bit pattern, e.g., an address pattern, and filter out received command channel information appropriately.

As described, one aspect of the commands sent by the processor into the receive circuit is that the commands can be canceled. An example of such a command is a speculative read operation, which may be canceled if the read operation is never actually executed by the code operating on the processor.

5 Referring to Fig. 5, operations of the logic circuit in Fig. 2 are illustrated in response to cancellation of a command. The command segments are unloaded from buffers 203-206 in 601 and stored into command queue 212 in 603. However, there is a possibility that the command will be canceled. As each of the command segments are stored in FIFOs 213-219, write pointers associated with the FIFOs, which indicate
10 where to write (also referred to herein as a push operation) the next command segment are incremented. Read pointers are also utilized to indicate the location from which to read (or pop) the next command. The push operations originate on the receiver side and the pop operations unload data from command queue 212 to the command interpreter.

15 Assume it takes four bit times to load the buffers 203-206 with a complete command. After the first bit time (or phase 1 of the command channel cycle), buffer 203, storing the first command segment, is unloaded into FIFO 213. After the second bit time, the second command segment, stored in buffer 204 is unloaded into FIFO 215. After the third bit time, the third command segment, stored in buffer 205 is
20 unloaded into FIFO 217. Assuming the cancellation indication is contained in the command segment received in the last bit time (or phase 4), the circuit in Fig. 2 determines if the current command segment is the last command segment in 507. If so, it is determined if a cancellation indication is asserted in 511. In one embodiment, one of the bits associated with the command segment received in the fourth bit time
25 for a read operation is a read valid bit. The read valid bit indicates whether the current read operation being loaded into command queue 212 is valid. Note that at the time the bit is checked for validity, data from the other three phases have already been loaded into command queue 212. If the bit indicates that the read operation is valid, then the circuit returns to unload a next command in 501.

30 If however the read valid bit, or other cancellation indication, indicates that the current command should be canceled in operation 511, then to prevent the command

interpreter from having to throw out canceled commands, the canceled command needs to be removed from command queue 212. In one implementation, that is accomplished by decrementing write pointers associated with each of the FIFOs storing segments of the canceled command in 513. A state machine can be used to
5 decrement each of the FIFO write pointers so the next push operation will place data into the same FIFO locations, which were used for the canceled read operations. Thus, in 515, it is determined if all write pointers have been decremented and if not, the control logic returns to 513 to decrement the next write pointer. The write pointers may be decremented in consecutive clock cycles. Once all write pointers
10 have been decremented, the control logic has completed the undo-push operation associated with that cancelled command and the next command segment is unloaded from the buffers in 501.

While the logic has detected that the last command segment includes a bit indicating the command is invalid, in order to keep the control logic for all of the
15 FIFOs in command queue 212 the same, the fourth command segment is stored into FIFO 219. That is illustrated at 509. While the operation of storing the last command segment in the command queue is shown as being performed before checking for the cancellation indication, in fact, the push operation may be performed after the check is performed in 511. It is shown prior to the check operation in 511 to illustrate that
20 in this particular implementation, the last segment is always stored in command queue 212. The state machine finishes by decrementing the write pointer for FIFO 219. After that is complete the write pointers for all the FIFOs point to the storage locations for the canceled command. In other implementations, only FIFOs already storing command segments of a canceled command need to have their write pointers
25 decremented and the last FIFO may never be written with the last command segment from a cancelled command. In either case, the push operation for the cancelled command has been undone.

In addition to maintaining a read and write pointer, a content count may be maintained, which is a counter that keeps track of the number of commands in the
30 command queue. If the content counter indicates that the command queue is full and another command arrives from the processor, then an error flag can be asserted, in some implementations, to indicate the error condition. When a command is pushed

into the queue, the content counter is incremented and when a command is popped from the queue the content counter is decremented. If the content counter indicates that the command queue is full and another command arrives to be pushed into the command queue, then an error flag can be asserted to indicate the error condition.

- 5 The content counter may be incremented as a result of an invalid command. If so, content counter is decremented as a result of a cancelled command.

In order to ensure the undo-push operation does not overlap with a push signal coming from the receive circuit, in one embodiment, the undo-push operation starts within one clock of receipt of the read valid bit indicating the command is canceled.

- 10 That is, the undo-push operation starts by the next bit time 0.

However, in an embodiment in which the undo-push operation does overlap with a push signal pushing a next command into the command queue, then the content counter may be adjusted as shown in Fig. 6, which illustrates control for the content counter. If a push operation 603 occurs by itself, the content counter is incremented by 1 in 609 and the write counter also is incremented. That is a normal push operation. If a push operation 603 and an undo-push operation 607 overlap, the content counter and write counter remain the same, and the circuit returns to idle state 601. If a push operation 603 and a pop operation 605 and an undo push operation 611 overlap, then the content counter is decremented by 1 in 621 and the write counter remains the same. If an undo-push operation 619 occurs alone or a pop operation 513 occurs alone, the content counter is decremented by 1 in 621. In the event of the undo-push operation occurring alone, then the write counter decrements. If a pop operation 613 and an undo-push operation 615 overlap, then the content counter is decremented by 2 in 617 and the write counter is decremented by 1.

- 25 The description of the invention set forth herein is illustrative, and is not intended to limit the scope of the invention as set forth in the following claims. For instance, while the invention has been described in an embodiment in which the receive circuit is located on a north bridge, one of skill in the art will appreciate that the teachings herein can be utilized any place it is desirable to provide undo-push capability, for example, on a processor with an integrated memory control function.
- 30 Variations and modifications of the embodiments disclosed herein, may be made

based on the description set forth herein, without departing from the scope and spirit of the invention as set forth in the following claims.